AD-A008 238

MODELING CONSIDERATIONS IN COMPUTER COMMUNICATION RESOURCE CONTROL

Stephen R. Kimbleton

University of Southern California

Prepared for:

Office of Naval Research
Advanced Research Projects Agency

January 1975

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 115124	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) MODELING CONSIDERATIONS IN COMPUTER COMMUNICATION RESOURCE CONTROL		5. TYPE OF REPORT & PERIOD COVERED Research Paper
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Stephen R. Kimbleton		8. CONTRACT OR GRANT NUMBER(s) N00014-67-A-0181-0036 and DAHC 15 72 C 0308
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of Southern California Information Sciences Institute 4676 Admiralty Way, Marina del Rey, California 90291		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ONR #NR 049-311 ARPA Order #2223/1
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Information Systems Program (OVER) Arlington, Virginia 22217		12. REPORT DATE January 1975
		13. NUMBER OF PAGES 16
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report)
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

This document approved for public release and sale; distribution is unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

Presented at the Eighth Hawaii International Conference on System Sciences, Honolulu, Hawaii, January 7-9, 1975.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Networking, network control, network performance, performance prediction, host sizing

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Effective utilization of mission oriented networks requires coordination and control of both communication and computer facilities. Achievement of these capabilities is shown to require a means for prediction of the effects of network modifications upon network performance. It is observed that standard monolithic simulation approaches are unlikely to provide an effective means for network performance prediction and the feasibility of developing a suitable structured analytically driven approach is explored.

11. Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

ACCESSION for

NTIS	White Section	☐
DDC	Buff Section	☐
UNANNOUNCED		☐

JUSTIFICATION ....

....

BY ....

DISTRIBUTION/AVAILABILITY CODES

Dist.	AVAIL. and/or SPECIAL	
A		

MODELING CONSIDERATIONS IN COMPUTER COMMUNICATION RESOURCE CONTROL

by

Stephen R. Kimbleton
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, California 90291

Abstract -- Effective utilization of mission oriented networks requires coordination and control of both communication and computer facilities. Achievement of these capabilities is shown to require a means for prediction of the effects of network modifications upon network performance. It is observed that standard monolithic simulation approaches are unlikely to provide an effective means for network performance prediction and the feasibility of developing a suitable structured analytically driven approach is explored.

Preparation of this paper was supported by the Advanced Research Projects Agency under Contract No. DAHC15 72 C 0308, ARPA Order No. 2223/1, Program Code No. 3D30 and 3P10, and the Office of Naval Research, Information Systems Program under Contract N00014-67-A-0181-0036 (NR 049-311).

The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed implied, of the Advanced Research Projects Agency, the Office of Naval Res ch, or the U.S. Government.

INTRODUCTION

At present, three major types of networks can be distinguished: (1) Terminal Access Networks, (2) Value Added Networks, and (3) Mission Oriented Networks. The first category is exemplified by some of the larger time-sharing services (e.g., TYMSHARE), which serve primarily to connect the user to physically remote resources. Value added networks (e.g., ARPANET) may be viewed as a logical extension of the concept of a computer utility; sharing of both hardware and software resources is possible, but the precise manner of sharing depends upon agreements between the parties concerned. An important characteristic of value added networks is the independence (both structural and organizational) of hosts and communication facilities--which, in the case of ARPANET, are referred to as the subnet.

By definition, a mission oriented network consists of a collection of computer and communication resources whose common objective is to achieve some organizational information processing function. Two major examples of such networks are the World Wide Military Command and Control System (WWMCCS) and the Air Force Advanced Logistics System (ALS).

The basic thesis of this paper is that effective utilization of mission oriented networks requires coordination and control of both communication and computer facilities, that network dynamics require these capabilities to be effectively 'on-line,' and that their achievement requires a model-based approach. The remainder of the paper outlines some of the implications of this thesis and discusses its feasibility.

CONTROL PROCEDURE COMPONENTS

Implementing a control procedure requires identification of (1) objectives, (2) decision space, and (3) the decision support system [SCOTM 74] used in selecting state variable based controls to achieve the objectives. In the complex, dynamic environment implicit in mission oriented networking, identifying these components will require an iterative approach. However, the current state of the art in computer system control provides a point of departure.

Control Objectives

The global objectives of network control are assumed to be the assurance of (1) acceptable workload processing through control of response/turnaround times, (2) adequate reliability, and (3) suitable job and data migration capabilities. Clearly these three objectives are interrelated. Thus, reliability may be achieved both through traditional static approaches of equipment redundancy or through dynamic approaches based on workload migration, the feasibility of acceptably processing a given workload is affected by the migration policies used, and the job and data migration capabilities desired are dependent upon network dynamics induced by workload processing requirements as well as by reliability requirements.

Control Decisions

Control decisions can be subdivided into organizational and technological decisions. Table 1 provides a partial list of some issues in each category. The distinction between the categories is not airtight; for instance, operator effectiveness levels (an organizational issue) may be strongly affected by the design and quality of technological control capabilities. The division between organizational and technological issues can be understood by observing that for organizational issues (1) management satisfices rather than optimizes, and (2) the determination of reasonable values is environmentally dependent. For technological issues, by contrast, the concept of an 'optimal' solution is appropriate, though elusive.

Although control of organizational issues is of obvious importance, its context-dependent nature makes dubious its amenability to centralized approaches [KIMBS 73]. Thus, the remainder of this paper is concerned with the control of technological issues. The importance of even this reduced domain of decisions is clear; moreover, if a solution methodology can be obtained, it can be used to determine the technological implications of organizational alternatives. Presenting such information to management would permit effective decisionmaking in spite of the general lack of well-defined utility functions for such issues [GROCJ 72].

TABLE 1

ISSUES IN COMPUTER SYSTEMS PERFORMANCE

ORGANIZATIONAL ISSUES	TECHNOLOGICAL ISSUES
Unnecessary jobs	Extended instruction set
Programmer Training	More host memory
Priority structure requirements	Device/channel ratio
Programming language support	Data set location

Technological Control Decisions

Technological control decisions can be structured into two categories: (1) parameter decisions and (2) host interaction decisions. Parameter decisions can be considered at four distinct levels: (1) the vendor, (2) the organizational information processing function, (3) the individual host, and (4) the individual user. The spectrum of available alternatives decreases as one passes downward through this list. For networks, the distinction between the vendor and the organizational information processing function has been somewhat nebulous, since (as is true for ARPANET) equipment design and usage is effectively controlled by the same agency. For the purposes of this discussion it is sufficient to assume that the organizational information processing function is concerned with the selection and configuration of essentially 'off the shelf' equipment. This observation appears to be essentially accurate for non-ARPA agencies intent upon using ARPANET technology.

Parameter decisions for an organization establishing an ARPANET-type computer communication system can be structured by considering the basic resource categories defining a (mission oriented) network: (1) transmission capabilities, (2) IMPs, and (3) host computers. Major components of the first category are: mode (e.g., satellite, land line, etc.), bandwidth, delay and error rate. Current IMP alternatives are limited and are reflected primarily in the choice of minicomputer used to drive the IMP [HEARF 70]. However, an essentially different IMP architecture permitting a broader spectrum of alternatives is becoming available [HEARF 73].

The possible control variables available for an individual host computer appear enormous at first sight. However, examination of management implemented control procedures reveals that these control variables can be structured into four categories: (1) job schedules, (2) hardware configurations, (3) software components, and (4) file device assignments. The preceding comments imply that, although the space of parameter decisions is large, it is not unbounded; further, it can be structured relatively tightly.

Interaction decisions affect both individual host operation and host cooperation. Although optimization of host operation is relatively well understood at a conceptual level [MILLE 73], the possibilities implicit in host cooperation for mission oriented networks do not yet appear to have received substantial study. Further, important tools such as Network Operating Systems, Network Job Control Languages, and other required protocols appear to be in the 'study' rather than implementation or utilization stage. New versions of RSEXEC [THOMR 73] are partially concerned with relevant issues.

The importance of this topic stems from the fact that network management differs from individual host management due to the possibility of migration of either jobs or data. This possibility, in turn, implies the feasibility of (1) execution of jobs originating from remote sites to use temporary excess capacity, (2) sizing for the average workload and using the network to handle peak processing requirements, and (3) achieving increased reliability through dynamic rather than static approaches, i.e., through job/file migration rather than acquisition of additional, usually underutilized, hardware. Achievement of effective migration capabilities will permit savings in both financial and human resources.

DECISION SUPPORT FOR TECHNOLOGICAL CONTROL

The ultimate objective of network control should be a self controlling network. In the absence of this capability, and in order to develop a basis for research toward its achievement, it is appropriate to develop a framework, i.e., decision support system, enabling effective man-machine control. Consideration of this problem [KIMBS 74A] reveals a requirement for two primary capabilities: (1) a report/display capability to identify outages, track evolving workload requirements, and serve as the basic information source to assist the brokerage function required to achieve effective workload migration, and (2) a performance projection mechanism for evaluation of the effects of management alternatives.

Although significant work is required to determine the 'right' combination of reports and the appropriate display mechanisms, current computer system control literature clearly reveals the feasibility of this objective. Nevertheless, significant functional differences will accrue due to the requirement that information be maintained and displayed 'on-line' to accommodate network dynamics.

On-line information is also required to permit useful projection of the effects of possible alternatives. Moreover, standard computer system simulation techniques appear to be unpromising due to their relatively slow execution rate; thus, it is usual for a computer system simulator providing information on device utilizations and delays to execute no faster than 1 to 10 times faster than real time. However, in a reasonably complex network environment, this implies an execution speed for a network simulator which is slower than real time and therefore unacceptable. Consequently, the remainder of this paper will be devoted to consideration of alternative means of network performance projection.

NETWORK PERFORMANCE PROJECTION

Performance projection may be attempted through either analytic or simulation-based approaches. Further, for a given approach either a monolithic or structured design may be used. Traditionally, simulation-based approaches tend to use a monolithic design while analytically based approaches use a carefully structured design which, additionally, usually incorporates a substantial (some would say excessive) level of simplifying assumptions.

A major advantage of the monolithic approach is its initial fidelity to the process (network) being modeled; an unfortunate corollary is the extreme level of complexity which may result with an attendant difficulty in distinguishing causes, effects and interactions. Further, the time required for this approach is sufficiently large that in one case completion of the simulator was roughly coincidental with completion of the project [NIELN 69]. Perhaps as a result, analytic approaches have enjoyed a continuing, and apparently expanding, success in modeling individual computer systems. Indeed, at a recent conference, four papers concerned with the development of analytic system models were presented [BARDY 74], [BROWJ 74], [KIMBS 74B] and [WYATJ 74]. In general, these approaches permit a substantially expanded level of detail by incorporating a hierarchical structure. We now comment upon the requirements implicit in the development of an analytically driven network simulator through usage of analytically driven host simulators.

The assumption that network performance can be predicted as a function of host, IMP and transmission line performance has not been completely validated and careful controlled testing is desirable to determine the extent to which this assumption is warranted [BELLT 73]. However, Kleinrock [KLEIL 64] has investigated a near relative of this assumption in the context of message switching networks. The results were generally affirmative. Further, this assumption is consistent with the results of previous work which has been done in the area of computer system modeling and which suggests that this assumption is generally warranted provided the system being modeled is operating under subsaturation conditions. However, since subsystem saturation tends to introduce a very unstable system behavior, performance prediction under conditions of subsystem saturation is likely to be unrewarding. Instead, detection and notification of the presence of saturation conditions should be investigated.

An analytically driven approach to prediction of host performance for production batch computer systems is described in [KIMBS 74B]. In general, the approach used is that of variable precision analytically driven models. Since this approach permits prediction of system performance as a function of synthetic module-like [SREEK 74] characterizations of the jobs comprising the system workload, its utilization as the host performance prediction mechanism of an analytically driven network simulator seems feasible. Indeed, the objective in developing this model was to obtain output information comparable with that provided by simulators, i.e. device utilizations and delays for jobs, shifts and time segments, with an execution speed of approximately two orders of magnitude faster than real time. Application of this approach to a network environment requires its extension to include interactive systems as well as a refinement of the approach used to represent channel effects.

Since an IMP is, in effect, a dedicated minicomputer with very specialized capabilities, it would seem feasible to seek to predict IMP performance using either simulation, analytically driven approaches described above, or a tabular technique. However, some care is required since prediction of IMP performance essentially involves characterization of program behavior, a topic not known for its tractability. Because of the importance of queuing delays in predicting IMP performance, explicit representation of some internal queues may prove appropriate. Thus, the approaches described in [BROWJ 74] and [BARDY 74] should prove of use.

Projection of transmission line performance requires some care. From a mathematical viewpoint, this problem would seem to be closely related to the problem of projection of the performance impact of data paths within a computer system. Although a limited amount of work concerned with this subject has appeared, careful validation seems to be required. It should be noted, however, that in the absence of such a model it is still feasible to perform controlled tests to establish the performance capabilities and limitations of a model in which all transmission lines are assumed to have essentially infinite full duplex bandwidth.

In summary, analytically driven projection of network performance appears feasible. Further, the speed achievable through this approach is required if management is to be able to rapidly predict the effects of alternative network configurations required for maintenance of appropriate traffic levels in the face of outages [KIMBS 74A].

REFERENCES

BARDY 74 Bard, Y., "An Analytic Model for CP-67 and VM/370," Preprint for IRIA Workshop on Computer Architectures and Networks, Rocquencourt, France, August 12-14, 1974.

BELLT 73 Bell, T. E., B. W. Boehm and R. A. Watson, COMPUTER PERFORMANCE ANALYSIS: FRAMEWORK AND INITIAL PHASES FOR A PERFORMANCE IMPROVEMENT EFFORT, The Rand Corporation, R-549-1-PR, November 1972.

BROWJ 74 Browne, J. C., K. M. Chandy, R. M. Brown, T. W. Keller, and

D. Towsley, "Hierarchical Techniques for Development of Realistic Models of Complex Computer Systems," Preprint for IRIA Workshop on Computer Architectures and Networks, Rocquencourt, France, August 12-14, 1974.

GROCJ 72 Grochow, J. M., TIME-SHARING CHARACTERISTICS AND USER UTILITY, Working Paper No. 582-72, Massachusetts Institute of Technology, Cambridge, Massachusetts, January, 1972.

HEARF 70 Heart, F. E., R. E. Kahn, S. M. Ornstein, W. R. Crowther, D. C. Walden, "The Interface Message Processor for the ARPA Computer Network," PROCEEDINGS OF AFIPS SJCC, 1970, Vol. 36.

HEARF 73 Heart, F. E., S. M. Ornstein, W. R. Crowther, and B. Barker, "A New Minicomputer/Multiprocessor for the ARPA Network," AFIPS CONFERENCE PROCEEDINGS, Vol. 42, 1973, pp. 529-537.

KIMBS 73 Kimbleton, S. R., "An Analytic Framework for Computer System Sizing and Tuning," The Rand Corporation, P-5162, January 1974. Also to appear in the Proceedings of the NBS/ACM Symposium on Performance Evaluation, San Diego, California, March 1973.

KIMBS 74A --- "Network Management Information Center," A RESEARCH PROGRAM IN THE FIELD OF COMPUTER TECHNOLOGY, ANNUAL TECHNICAL REPORT: MAY 1973--MAY 1974, University of Southern California, Information Sciences Institute, ISI/SR-74-2, pp. 97-106.

KIMBS 74B --- "A Fast Approach to Computer System Performance Prediction," Preprint for the IRIA Workshop on Computer Architectures and Networks, Rocquencourt, France, E. Gelenbe and R. Mahl (Eds.), North-Holland Publishing Company, Inc., 1974, pp. 251-262.

KLEIL 64 Kleinrock, L., COMMUNICATION NETS: STOCHASTIC MESSAGE FLOW AND DELAY, McGraw-Hill, New York, 1964.

MILLE 73 Miller, E. F., "Bibliography and KWIC Index on Computer Performance Measurement," General Research Corporation, RM-1809, June 1973.

NIELN 67 Nielsen, N. R., "Computer Simulation of Computer System Performance," PROCEEDINGS ACM 22ND NATIONAL CONFERENCE, August 1967, pp. 581-590.

SCOTM 74 Scott Morton, M. S., "Decision Support Systems: Some Lessons From an On-Going Application," IFIP CONGRESS 74 PROCEEDINGS, Stockholm, Sweden, August 5-10, 1974, pp. 1062-1066.

SREEK 74 Sreenivasan, K. and A. J. Kleinman, "On the Construction of a Representative Synthetic Workload," COMMUNICATIONS OF THE ACM, Vol. 27, No. 3, 1974, pp. 127-132.

THOMR 73 Thomas, R. H., "A Resource Sharing Executive for the ARPANET," AFIPS CONFERENCE PROCEEDINGS, Vol. 42, 1973, pp. 155-163.

WYATJ 74 Wyatt, J. B., "Evaluating Alternatives in Information System Structure by Simulation," Preprint for the IRIA Workshop on Computer Architectures and Networks, E. Gelenbe and R. Mahl, (Eds.), August 12-14, 1974, Rocquencourt, France, pp. 405-417.